

Faster Algorithms for Semi-Matching Problems *

Jittat Fakcharoenphol
Kasetsart University
jittat@gmail.com

Bundit Laekhanukit[†]
McGill University
blaekh@cs.mcgill.ca

Danupon Nanongkai[†]
University of Vienna
danupon@gmail.com

Abstract

We consider the problem of finding *semi-matching* in bipartite graphs which is also extensively studied under various names in the scheduling literature. We give faster algorithms for both weighted and unweighted cases.

For the weighted case, we give an $O(nm \log n)$ -time algorithm, where n is the number of vertices and m is the number of edges, by exploiting the geometric structure of the problem. This improves the classical $O(n^3)$ -time algorithms by Horn [Operations Research 1973] and Bruno, Coffman and Sethi [Communications of the ACM 1974].

For the unweighted case, the bound could be improved even further. We give a simple divide-and-conquer algorithm which runs in $O(\sqrt{nm} \log n)$ time, improving two previous $O(nm)$ -time algorithms by Abraham [MSc thesis, University of Glasgow 2003] and Harvey, Ladner, Lovász and Tamir [WADS 2003 and Journal of Algorithms 2006]. We also extend this algorithm to solve the *Balance Edge Cover* problem in $O(\sqrt{nm} \log n)$ time, improving the previous $O(nm)$ -time algorithm by Harada, Ono, Sadakane and Yamashita [ISAAC 2008].

Categories and Subject Descriptors: G.2.2 [Graph Theory]–Graph algorithms

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Semi-Matching, Scheduling, Combinatorial Optimization, Design and Analysis of Algorithm

1 Introduction

In this paper, we consider a relaxation of the maximum bipartite matching problem called *semi-matching* problem, in both weighted and unweighted cases. This problem has been previously studied in the scheduling literature under different names, mostly known as (non-preemptive) scheduling independent jobs on unrelated machines to minimize flow time, or $R||\sum C_j$ in the standard scheduling notation [3, 26, 2].

Informally, the problem can be explained by the following off-line load balancing scenario: We are given a set of jobs and a set of machines. Each machine can process one job at a time and it takes different amounts of time to process different jobs. Each job also requires different processing times if processed by different machines. One natural goal is to have all jobs processed with the

*A preliminary version of this paper appeared in ICALP'10 [11].

[†]Most of the work was done while all authors were at Kasetsart University.

minimum *total completion time*, or *total flow time*, which is the summation of the duration each job has to wait until it is finished. Observe that if the assignment is known, the order each machine processes its assigned jobs is clear: It processes jobs in an increasing order of the processing time.

To be precise, the semi-matching problem is as follows. Let $G = (U \cup V, E)$ be a weighted bipartite graph, where U is a set of jobs and V is a set of machines. For any edge uv , let w_{uv} be its weight. Each weight of an edge uv indicates time it takes v to process u . Through out this paper, let n denote the number of vertices and m denote the number of edges in G . A set $M \subseteq E$ is a *semi-matching* if each job $u \in U$ is incident with exactly one edge in M . For any semi-matching M , we define the *cost* of M , denoted by $\text{cost}(M)$, as follows. First, for any machine $v \in V$, its cost with respect to a semi-matching M is

$$\text{cost}_M(v) = (w_1) + (w_1 + w_2) + \dots + (w_1 + \dots + w_{\deg_M(v)}) = \sum_{i=1}^{\deg_M(v)} (\deg_M(v) - i + 1) \cdot w_i$$

where $\deg_M(v)$ is the degree of v in M and $w_1 \leq w_2 \leq \dots \leq w_{\deg_M(v)}$ are weights of the edges in M incident with v sorted increasingly. Intuitively, this is the total completion time of jobs assigned to v . Note that for the unweighted case (i.e., when $w_e = 1$ for every edge e), the cost of a machine v is simply $\deg_M(v) \cdot (\deg_M(v) + 1)/2$. Now, the cost of the semi-matching M is simply the summation of the cost over all machines:

$$\text{cost}(M) = \sum_{v \in V} \text{cost}_M(v).$$

The goal is to find an *optimal semi-matching*, a semi-matching with minimum cost.

Previous works: Although the name “semi-matching” was recently proposed by Harvey, Ladner, Lovász, and Tamir [19], the problem was studied as early as 1970s when an $O(n^3)$ algorithm was independently developed by Horn in [20] and by Bruno, Coffman and Sethi in [6]. Since then no progress has been made on this problem except on its special cases and variations. For the special case of *inclusive set restriction* where, for each pair of jobs u_1 and u_2 , either all neighbors of u_1 are neighbors of u_2 or vice versa, a faster algorithm with $O(n^2)$ running time was given by Spyropoulos and Evans [40]. Many variations of this problem were proved to be NP-hard, including the preemptive version [39], the case when there are deadlines [41], and the case of optimizing total weighted tardiness [29]. The variation where the objective is to minimize $\max_{v \in V} \text{cost}_M(v)$ was also considered [32, 25].

The unweighted case of the semi-matching problem also received considerably attention in the past few years. Since it was shown by [19] that an optimal solution of the semi-matching problem is also optimal for the makespan version of the scheduling problem (where one wants to minimize the time the last machine finishes), we mention the results of both problems. The problem was first studied in a special case, called *nested* case where, for any two jobs, if their sets of neighbors are not disjoint, then one of these sets contains the other set. This case is shown to be solvable in $O(m + n \log n)$ time [36, p.103]. For the general unweighted semi-matching problem, Abraham [1, Section 4.3] and Harvey, Ladner, Lovász and Tamir [19] independently developed two algorithms with $O(nm)$ running time. Lin and Li [28] also gave an $O(n^3 \log n)$ -time algorithm which is later generalized to a more general cost function [27]. Recently, Lee, Leung and Pinedo [24] showed that the problem can be solved in polynomial time even when there are release times.

The unweighted semi-matching problem is recently generalized to the quasi-matching problem by Bokal, Bresar and Jerebic [4]. In this problem, a function g is provided and each vertex $u \in U$

is required to connect to at least $g(u)$ vertices in v . Therefore, the semi-matching problem is when $g(u) = 1$ for every $u \in U$. They also developed an algorithm for this problem which is a generalization of the Hungarian method and used it to deal with a routing problem in CDMA-based wireless sensor networks.

Motivated by the problem of assigning wireless stations (users) to access points, the unweighted semi-matching problem is also generalized to the problem of finding optimal semi-matching with minimum weight where an $O(n^2m)$ time algorithm is given [15].

Approximation algorithms and online algorithms for this problem (both weighted and unweighted cases) and the makespan version have also gained a lot of attention over the past few decades and have applications ranging from scheduling in hospital to wireless communication network. (See [26, 48] for the recent surveys.)

Applications: As motivated by Harvey et al. [19], even in an online setting where jobs arrive and depart over time, they may be reassigned from one machine to another cheaply if the algorithm's runtime is significantly faster than the arrival/departure rate. (One example of such case is the Microsoft Active Directory system [14, 19].) The problem also arose from the Video on Demand (VoD) systems where the load of video disks needs to be balanced while data blocks from the disks are retrieved or while serving clients [31, 45]. The problem, if solved in the distributed setting, can be used to construct a load balanced data gathering tree in sensor networks [37, 33]. The same problem also arose in peer-to-peer systems [42, 23, 43].

In this paper, we also consider an “edge cover” version of the problem. In some applications such as sensor networks, there are no jobs and machines but the sensor nodes have to be clustered and each cluster has to pick its own head node to gather information from other nodes in the cluster. Motivated by this, Harada, Ono, Sadakane and Yamashita [16] introduced the *balanced edge cover* problem¹ where the goal is to find an edge cover (set of edges incident to every vertex) that minimizes the total cost over all vertices. (The cost on each vertex is as previously defined.) They gave an $O(nm)$ algorithm for this problem and claimed that it could be used to solve the semi-matching problem as well. We show that this problem can be efficiently reduced to the semi-matching problem. Thus, our algorithm (for unweighted case) also gives a better bound on the balanced edge cover problem.

Our results and techniques

We consider the semi-matching problem and give a faster algorithm for each of the weighted and unweighted cases. We also extend the algorithm for the unweighted case to solve the balanced edge cover problem.

- **Weighted Semi-Matching:** (Section 2) We present an $O(nm \log n)$ algorithm, improving the previous $O(n^3)$ algorithm by Horn [20] and Bruno et al. [6]. As in the previous results [20, 5, 17], we use the reduction of the weighted semi-matching problem to the weighted bipartite matching problem as a starting point. We, however, only use the structural properties arising from the reduction and do not actually perform the reduction.
- **Unweighted Semi-Matching:** (Section 3) We give an $O(\sqrt{nm} \log n)$ algorithm, improving

¹This problem is also known as a *constant jump system* (see, e.g., [44, 30]).

the previous $O(nm)$ algorithms by Abraham [1] and Harvey et al. [19].² Our algorithm uses the same reduction to the min-cost flow problem as in [19]. However, instead of cancelling one negative cycle in each iteration, our algorithm exploits the structure of the graphs and the cost functions to cancel many negative cycles in a single iteration. This technique can also be generalized to any convex cost function.

- **Balanced Edge Cover:** (Section 4) We also present a reduction from the balanced edge cover problem to the unweighted semi-matching problem. This leads to an $O(\sqrt{nm} \log n)$ algorithm for the problem, improving the previous $O(nm)$ algorithm by Harada et al. [16]. The main idea is to identify the “center” vertices of all the clusters in the optimal solution. (Note that any balanced edge cover (in fact, any minimal edge cover) clusters the vertices into stars.) Then, we partition the vertices into two sides, center and non-center ones, and apply the semi-matching algorithm on this graph.

2 Weighted semi-matching

In this section, we present an algorithm that finds optimal weighted semi-matching in $O(nm \log n)$ time.

Overview

Our improvement follows from studying the reduction from the weighted semi-matching problem to the weighted bipartite matching problem considered in the previous works [20, 6, 17] and the Edmonds-Karp-Tomizawa (EKT) algorithm for finding the weighted bipartite matching [9, 47]. We first review these briefly. For more detail, see Appendix A and B.

Reduction: As in [20, 6, 17], we consider the reduction from the semi-matching problem on bipartite graph $G = (U \cup V, E)$ to the minimum-weight bipartite matching on a graph \hat{G} . The reduction is done by exploding the vertices in V , i.e., for each vertex $v \in V$ we create $\deg(v)$ vertices, $v^1, v^2, \dots, v^{\deg(v)}$. We also make copies of edges incident to v in the original graph G , i.e., for each vertex $u \in U$ such that $uv \in E$, we create edges $uv^1, uv^2, \dots, uv^{\deg(v)}$. For each edge uv^i incident to v^i in \hat{G} , we set its weight to i times its original weight in G , i.e., $w_{uv^i} = i \cdot w_{uv}$. We denote the set of these vertices by \hat{V}_v . Thus, we have

$$\begin{aligned}\hat{G} &= (U \cup \hat{V}, \hat{E}) \\ \hat{V} &= \{v^1, v^2, \dots, v^{\deg_G(v)} : v \in V\} \\ \hat{E} &= \{uv^1, uv^2, \dots, uv^{\deg_G(v)} : uv \in E\} \\ \hat{w}_{uv^i} &= i \cdot w_{uv} \quad \forall uv \in E, i \in \{1, 2, \dots, \deg_G(v)\}\end{aligned}$$

The correctness of this reduction can be seen by replacing the edges incident to v in the semi-matching by the edges incident to v^1, v^2, \dots with weights in an decreasing order. For example, in Figure 1(a), edge u_1v_1 and edge u_2v_1 in the semi-matching in G correspond to $u_1v_1^1$ and $u_2v_1^2$ in the matching in \hat{G} . The reduction is illustrated in Figure 1(a).

²We also observe an $O(n^{5/2} \log n)$ algorithm that arises directly from the reduction by applying [21].




Figure 1:

This alone does not give an improvement on the semi-matching problem because the number of edges becomes $O(nm)$. However, we can apply some tricks to improve the running time. (See Appendix B.)

EKT algorithm: Our improvement comes from studying the behavior of the EKT algorithm for finding the bipartite matching in \hat{G} . The EKT algorithm iteratively increases the cardinality of the matching by one by finding a shortest augmenting path. Such path can be found by applying Dijkstra's algorithm on the *residual graph* D_M (corresponding to a matching M) with a *reduced cost*, denoted by \tilde{w} as an edge length.

Figure 1(b) shows examples of a residual graph D_M . The direction of an edge depends on whether it is in the matching or not. The weight of each edge depends on its weight in the original graph and the costs on its end vertices. We draw an edge of length 0 from s to all vertices in U_M and from all vertices in \hat{V}_M to t , where U_M and \hat{V}_M are the sets of unmatched vertices in U and \hat{V} , respectively. We want to find the shortest path from s to t or, equivalently, from U_M to \hat{V}_M .

The reduced cost is computed from the *potentials* on the vertices, which can be found as in Algorithm 2.1.³

Applying EKT algorithm directly leads to an $O(n(n' \log n' + m'))$ where $n = |U|$, $n' = |U \cup V|$ and m' is the number of edges in \hat{G} . Since $n' = |\hat{V}| = \Theta(m)$ and $m' = \Theta(n^2)$, the running time is $O(nm \log n + n^3)$. (We note that this could be brought down to $O(n^3)$ by applying the result of Kao, Lam, Sung and Ting [21] to reduce the number of participating edges. See Appendix B.) The bottleneck here is the Dijkstra's algorithm which needs $O(n' \log n' + m')$ time. We now review this algorithm and pinpoint the part that will be sped up.

Dijkstra's algorithm: Recall that the Dijkstra's algorithm starts from a source vertex and keeps adding to its shortest path tree a vertex with minimum tentative distance. When a new vertex v is added, the algorithm updates the tentative distance of all vertices outside the tree by relaxing *all* edges incident to v . On an n' -vertex m' -edge graph, it takes $O(\log n')$ time (using priority queue) to find a new vertex to add to the tree and hence $O(n' \log n')$ in total. Further, relaxing all edges

³Note that we set the potentials in an unusual way: We keep potentials of the unmatched vertices in \hat{V} to 0. The reason is roughly that we can speed up the process of finding the distances of all vertices but vertices in \hat{V}_M . Notice that this type of potentials is valid too (i.e., \tilde{w} is non-negative) since for any edge uv such that $v \in \hat{V}_M$ is unmatched, $\tilde{w}_{uv} = w_{uv} + p(u) - p(v) = w_{uv} + p(u) \geq 0$.

Algorithm 2.1 EKT ALGORITHM (\hat{G}, w)

```
1: Let  $M = \emptyset$ .
2: For every node  $v$ , let  $p(v) = 0$ . ( $p(v)$  is a potential on  $v$ .)
3: repeat
4:   Let  $\tilde{w}_{uv} = w_{uv} + p(u) - p(v)$  for every edge  $uv$ . ( $\tilde{w}_{uv}$  is a reduced cost of an edge  $uv$ .)
5:   For every node  $v$ , compute the distance  $d(v)$  which is the distance from  $U_M$  (the set of
     unmatched vertices in  $U$ ) to  $v$  in  $D_M$ . (Recall that the length of edges in  $D_M$  is  $\tilde{w}$ .)
6:   Let  $P$  be the shortest  $U_M$ - $\hat{V}_M$  path in  $D_M$ .
7:   Update the potential  $p(u)$  to  $d(u)$  for every vertex  $u \in U \cup (\hat{V} \setminus \hat{V}_M)$ .
8:   Augment  $M$  along  $P$ , i.e.,  $M = P \triangle M$  (where  $\triangle$  denotes the symmetric difference operator).
9: until all vertices in  $U$  are matched
10: return  $M$ 
```

takes $O(m')$ time in total. Recall that in our case, $m' = \Theta(n^2)$ which is too large. *Thus, we wish to reduce the number of edge relaxations to improve the overall running time.*

Our approach: We reduce the number of edge relaxation as follows. Suppose that a vertex $u \in U$ is added to the shortest path tree. For every $v \in V$, a neighbor of u in G , we relax all edges uv^1, uv^2, \dots, uv^i in \hat{G} *at the same time*. In other words, instead of relaxing $\Theta(nm)$ edges in \hat{G} separately, we group the edges to m groups (according to the edges in G) and relax all edges in each group together. We develop a relaxation method that takes $O(\log n)$ time per group. In particular, we design a data structure H_v , for each vertex $v \in V$, that supports the following operations.

- **RELAX**(uv, H_v): This operation works as if it relaxes edges uv^1, uv^2, \dots
- **ACCESSMIN**(H_v): This operation returns a vertex v^i (exploded from v) with minimum tentative distance among vertices that are not deleted (by the next operation).
- **DELETEMIN**(H_v): This operation finds v^i from **ACCESSMIN** and then returns and deletes v^i .

Our main result is that, by exploiting the structure of the problem, one can design H_v that supports **RELAX**, **ACCESSMIN** and **DELETEMIN** in $O(\log n)$, $O(1)$ and $O(\log n)$ respectively. Before showing such result, we note that speeding up Dijkstra's algorithm and hence EKT algorithm is quite straightforward once we have H_v : We simply build a binary heap H whose nodes correspond to vertices in an original graph G . For each vertex $u \in U$, H keeps track of its tentative distance. For each vertex $v \in V$, H keeps track of its *minimum tentative distance* returned from H_v .

Main idea: Before going into details, we sketch the main idea here. The data structure H_v that allows fast “group relaxation” operation can be built because of the following nice structure of the reduction: For each edge uv of weight w_{uv} in G , the weights $w_{uv^1}, w_{uv^2}, \dots$ of the corresponding edges in \hat{G} increase linearly (i.e., $w_{uv}, 2w_{uv}, 3w_{uv}, \dots$). This enables us to know the order of vertices, among v^1, v^2, \dots , that will be added to the shortest path tree. For example, in Figure 1(b), when $M = \emptyset$, we know that, among v^1 and v^2 , v^1 will be added to the shortest path tree first as it always has a smaller tentative distance.

However, since the length of edges in D_M does not solely depend on the weights of the edges in \hat{G} (in particular, it also depends on a potentials on both end vertices), it is possible (after some iterations of the EKT algorithm) that v^1 is added to the shortest path tree after v^2 .

Fortunately, due to the way the potential is defined by the EKT algorithm, a similar nice property still holds: Among v^1, v^2, \dots in D_M corresponding to v in G , if a vertex v^k , for some k , is added to the shortest path tree first, then the vertices on each side of v^k have a nice order: Among v^1, v^2, \dots, v^{k-1} , the order of vertices added to the shortest path tree is $v^{k-1}, v^{k-2}, \dots, v^2, v^1$. Further, among v^{k+1}, v^{k+2}, \dots , the order of vertices added to the shortest path tree is v^{k+1}, v^{k+2}, \dots .

This main property, along with a few other observations, allow us to construct the data structure H_v . In the next section, we show the properties we need and use them to construct H_v in the latter section.

2.1 Properties of the tentative distance

Consider any iteration of the EKT algorithm (with a potential function p and a matching M). We study the following functions f_{*v} and g_{*v} .

Definition 2.1. For any edge uv from U to V and any integer $1 \leq i \leq \deg(v)$, let

$$g_{uv}(i) = d(u) + p(u) + i \cdot w_{uv} \quad \text{and} \quad f_{uv}(i) = g_{uv}(i) - p(v^i) = d(u) + p(u) - p(v^i) + i \cdot w_{uv}.$$

For any $v \in V$ and $i \in [\deg(v)]$, define the lower envelope of f_{uv} and g_{uv} over all $u \in U$ as

$$f_{*v}(i) = \min_{u: uv \in E} f_{uv}(i) \quad \text{and} \quad g_{*v}(i) = \min_{u: uv \in E} g_{uv}(i).$$

Our goal is to understand the structure of the function f_{*v} whose values $f_{*v}(1), f_{*v}(2), \dots$ are tentative distances of v^1, v^2, \dots , respectively. The function g_{*v} is simply f_{*v} with the potential of v ignored. We define g_{*v} as it is easier to keep track of since it is a combination of linear functions g_{uv} and therefore piecewise linear. Now we state the key properties that enable us to keep track of f_{*v} efficiently. Recall that v^1, v^2, \dots are the exploded vertices of v (from the reduction).

Proposition 2.2. Consider a matching M and a potential p at any iteration of the EKT algorithm.

- (1) For any vertex $v \in V$, there exists α_v such that v^1, \dots, v^{α_v} are all matched and $v^{\alpha_v+1}, \dots, v^{\deg(v)}$ are all unmatched.
- (2) For any vertex $v \in V$, g_{*v} is a piecewise linear function.
- (3) For any edge $uv \in E$ where $u \in U$ and $v \in V$, and any i , $f_{uv}(i) = f_{*v}(i)$ if and only if $g_{uv}(i) = g_{*v}(i)$.
- (4) For any edge $uv \in E$ where $u \in U$ and $v \in V$, let α_v be as in (1). There exists an integer $1 \leq \gamma_{uv} \leq \alpha_v$ such that for $i = 1, 2, \dots, \gamma_{uv}-1$, $f_{uv}(i) \geq f_{uv}(i+1)$ and for $i = \gamma_{uv}, \gamma_{uv}+1, \dots, \alpha_v-1$, $f_{uv}(i) \leq f_{uv}(i+1)$. In other words, $f_{uv}(1), f_{uv}(2), \dots, f_{uv}(\alpha_v)$ is a unimodal sequence.

Figure 2(a) and 2(b) show the structure of g_{*v} and f_{*v} according to statement (2) and (4) in the above proposition. By statement (3), the two pictures can be combined as in Figure 2(c): g_{*v} indicates u that makes both g_{*v} and f_{*v} minimum in each interval and one can find i that minimizes f_{*v} in each interval by looking at α_v (or near α_v in some case).

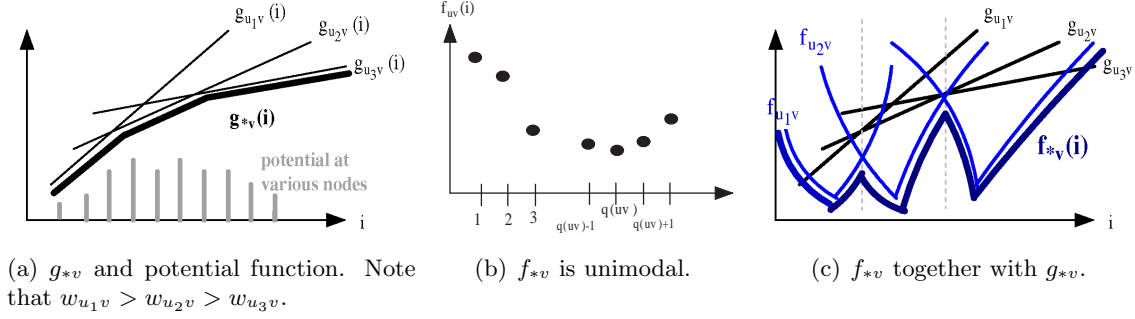


Figure 2:

Proof.

(1) The first statement follows from the following claim.

Claim 2.3. *For any i , if the exploded vertex v^{i+1} of v (in \hat{V}_v) is matched by M , then v^i is also matched.*

Proof. The claim follows from the fact that EKT algorithm maintains M so that M is an extreme matching. Suppose that v^{i+1} is matched by M (i.e., $uv^{i+1} \in M$), but v^i is not matched. Then we can remove uv^{i+1} from M and add uv^i to M . The resulting matching will have a cost less than M but have the same cardinality, a contradiction. \square

(2) To see the second statement, notice that $g_{uv} = d(u) + p(u) + i \cdot w_{uv}$ is linear for a fixed $uv \in E$. Hence, g_{*v} is a lower envelope of a linear function implying that it is piecewise linear.

(3) To prove the third statement, recall that for any u and any i , $f_{uv}(i) = g_{uv}(i) - p(v^i)$. Therefore, for any u, u' and i , $f_{uv}(i) > f_{u'v}(i)$ if and only if $g_{uv}(i) > g_{u'v}(i)$. Thus, the third statement follows.

(4) For the fourth claim, we first explain the intuition. First, observe that the function g_{uv} is increasing with rate w_{uv} . Moreover, the difference of $f_{uv}(i)$ and $f_{uv}(j)$ is a function of the potential $p(v^i)$ and $p(v^j)$ and the multiple of edge weight $(j-i)w_{uv}$. In fact, whether the difference is negative or positive depends on the value of these three parameters. We show that these parameters change monotonically and so we have the desired property.

To prove the fourth statement formally. We first prove two claims.

For the first claim below, recall that the potential of matched vertices, at any iteration, is defined to be the distance on the residual graph of the previous iteration. In particular, for any $v^i \in \hat{V}$, there is a vertex $u \in U$ such that $p(u) + i \cdot w_{uv} = p(v)$. (See Algorithm 2.1.)

Claim 2.4. *For any integer $i < \alpha_v$, consider the exploded vertices v^i and v^{i+1} . Let u and u' denote two vertices in U such that $p(u) + i \cdot w_{uv} = p(v^i)$ and $p(u') + (i+1) \cdot w_{u'v} = p(v^{i+1})$. Then $w_{uv} \geq p(v^{i+1}) - p(v^i) \geq w_{u'v}$.*

Proof. The first part, $w_{u'v} \geq p(v^{i+1}) - p(v^i)$, follows from $p(v^i) = p(u) + i \cdot w_{uv}$ and $p(v^{i+1}) \leq p(u) + (i+1) \cdot w_{uv}$. The second part, $p(v^{i+1}) - p(v^i) \geq w_{u'v}$, follows from, $p(v^i) \leq p(u') + i \cdot w_{u'v}$ and $p(v^{i+1}) = p(u') + (i+1) \cdot w_{u'v}$. \square

Proof of the next claim follows directly from the definition of f_{uv} (cf. Definition 2.1).

Claim 2.5. *For any $i < \alpha_v$, $f_{uv}(i) > f_{uv}(i+1)$ if and only if $p(v^{i+1}) - p(v^i) > w_{uv}$ and $f_{uv}(i) < f_{uv}(i+1)$ if and only if $p(v^{i+1}) - p(v^i) < w_{uv}$.*

Now, the fourth statement in the Proposition follows from the following statements: For any integer $i < \alpha_v$,

- (i) if $f_{uv}(i) > f_{uv}(i+1)$, then $f_{uv}(j) \geq f_{uv}(j+1)$ for any integer $j < i$, and
- (ii) if $f_{uv}(i) < f_{uv}(i+1)$, then $f_{uv}(j) \leq f_{uv}(j+1)$ for any integer $i \leq j \leq \alpha_v$.

To prove the first statement, let u' be such that $p(u') + i \cdot w_{u'v} = p(v^i)$. If $f_{uv}(i) > f_{uv}(i+1)$, then

$$p(v^i) - p(v^{i-1}) \geq w_{u'v} \geq p(v^{i+1}) - p(v^i) > w_{uv}$$

where the first two inequalities follow from Claim 2.4 and the third inequality follows from Claim 2.5. It then follows from Claim 2.5 that $f_{uv}(i-1) > f_{uv}(i)$. The first statement follows by repeating the argument above. The second statement can be proved similarly. This completes the proof of the fourth statement. \square

2.2 Data structure

Specification: Let us first redefine the problem so that we can talk about the data structure in a more general way. We show how to use this data structure for the semi-matching problem in the next section.

Let n and N be positive integers and, for any integer i , define $[i] = \{1, 2, \dots, i\}$. We would like to maintain at most n functions f_1, f_2, \dots, f_n mapping $[N]$ to a set of positive reals. We assume that f_i is given as an *oracle*, i.e., we can get $f_i(x)$ by sending a query x to f_i in $O(1)$ time.

Let L and S be a subset of $[N]$ and $[n]$, respectively. (As we will see shortly, we use L to keep the numbers left undeleted in the process and S to keep the functions inserted to the data structure.) Initially, $L = [N]$ and $S = \emptyset$. For any $x \in [N]$, let $f_S^*(x) = \min_{f_i \in S} f_i(x)$. We want to construct a data structure \mathcal{H} that supports the following operations.

- **ACCESSMIN(\mathcal{H}):** Return $x \in L$ with minimum value f_S^* , i.e., $x = \arg \min_{x \in L} f_S^*(x)$.
- **INSERT(f_i, \mathcal{H}):** Insert f_i to S .
- **DELETEMIN(\mathcal{H}):** Delete x from L where x is returned from ACCESSMIN(\mathcal{H}).

Properties: We assume that f_1, f_2, \dots have the following properties.

- For all i , f_i is *unimodal*, i.e., there is some $\gamma_i \in [N]$ such that $f_i(1) \geq f_i(2) \geq \dots \geq f_i(\gamma_i) \leq f_i(\gamma_i + 1) \leq f_i(\gamma_i + 2) \leq \dots \leq f_i(N)$. We assume that γ_i is given along with f_i .
- We also assume that each f_i comes along with a linear function g_i where, for any $x \in [N]$, $g_i(x) = x \cdot w_i + d_i$, for some w_i and d_i . These linear functions have a property that $f_i(x) = f_S^*(x)$ if and only if $g_i(x) = g_S^*(x)$, where $g_S^*(x) = \min_{g_i \in S} g_i(x)$.
- Finally, we assume that once x is deleted from L , $f_S^*(x)$ will never change, even after we add more functions to S .

For simplicity, we also assume that $w_i \neq w_j$ for all $i \neq j$. This assumption can be removed by taking care of the case of equal weight in the insert operation. We now show that there is a data structure such that every operation can be done in $O(\log n)$ time.

Data structure design: We have two data structures to maintain the information of f_i 's and g_i 's. First, we create a data structure T_g to maintain an ordered sequence g_{i_1}, g_{i_2}, \dots such that $w_{i_1} \geq w_{i_2} \geq \dots$. We want to be able to insert a new function g_i to T_g in $O(\log n)$ time. Moreover, for any w , we want to be able to find w_{i_j} and $w_{i_{j+1}}$ such that $w_{i_j} \leq w < w_{i_{j+1}}$ in $O(\log n)$ time. Such T_g can be implemented by a balanced binary search tree, e.g., an AVL tree.

Observe that the linear functions g_{i_1}, g_{i_2}, \dots appear in the lower envelope in order, i.e., if $g_{i_j}(x) \geq g_{i_{j+1}}(x)$, then $g_{i_j}(y) \geq g_{i_{j+1}}(y)$ for any $y > x$. Therefore, we can use data structure T_g to maintain the range of values such that each g_i (and therefore f_i) is in the lower envelope. That is, we use T_g to maintain $x_1 \leq y_1 \leq x_2 \leq y_2 \leq \dots$ such that $g_i(x) = g_S^*(x)$ for all i and $x_i \leq x \leq y_i$.

Consider the value $\min_{x \in \{x_i, x_i+1, \dots, y_i\} \cap L} f_i(x)$. Since f_i is unimodal, the minimum value of $f_i(x)$ over $\{x_i, x_i+1, \dots, y_i\} \cap L$ attains at the point closest to γ_i either from the left or from the right. Thus, we can use two pointers p_i and q_i such that $x_i \leq p_i \leq \gamma_i \leq q_i \leq y_i$ to maintain the minimum value of f_i from the left and right of γ_i , i.e., the minimum value $\min_{x \in \{x_i, x_i+1, \dots, y_i\} \cap L} f_i(x)$ is either $f_i(p_i)$ or $f_i(q_i)$. Finally, we use a binary heap B to store the values $f_1(p_1), f_2(p_2), \dots$ and $f_1(q_1), f_2(q_2), \dots$ so that we can search and delete the minimum among these values in $O(\log n)$ time.

More details of the implementation of each operation are the followings.

- **ACCESSMIN(\mathcal{H}):** This operation is done by returning the minimum value in B . This value is $\min(f_1(p_1), f_2(p_2), \dots, f_1(q_1), f_2(q_2), \dots) = \min_{x \in L} f_S^*(x)$.
- **INSERT(f_i, \mathcal{H}):** First, insert g_i to T_g which can be done as follows. Let the current ordered sequence be g_{i_1}, g_{i_2}, \dots . In $O(\log n)$ time, we find g_{i_j} and $g_{i_{j+1}}$ such that $w_{i_j} \leq w_i < w_{i_{j+1}}$ and insert g_i between them. Moreover, we update the region g_{i_j}, g_i , and $g_{i_{j+1}}$ are in the lower envelope of g_S^* , i.e., we get the values $y_{i_j}, x_i, y_i, x_{i_{j+1}}, y_{i_{j+1}}$ (note that $y_{i_j} \leq x_i \leq y_i \leq x_{i_{j+1}} \leq y_{i_{j+1}}$).

Next, we deal with the pointers p_i and q_i : We set $p_i = \min(\gamma_i, y_i)$ and $q_i = \max(\gamma_i, x_i)$. (The intuition here is that we would like to set $p_i = q_i = \gamma_i$ but it is possible that $\gamma_i < x_i$ or $\gamma_i > y_i$ which means that γ_i is not in the region that g_i is in the lower envelope g_S^*). Finally, we also update p_{i_j} and $q_{i_{j+1}}$: $p_{i_j} = \min(p_{i_j}, x_i)$ and $q_{i_{j+1}} = \max(q_{i_{j+1}}, y_i)$. Figure 3 shows an effect of inserting a new function.

We note one technical detail here: It is possible that p_i is already deleted from L . This implies that there is another function $f_{i_{j'}}$ such that $f_{i_{j'}}(p_i) = f_i(p_i)$ (since we assume that if p_i is already deleted, then $f_S^*(p_i)$ will never change even when we add more functions to S). There are two cases: $j' < j$ or $j' > j$. For the former case, we know that $f_{i_{j'}}(p_i - 1) < f_i(p_i - 1)$ since $w_{j'} > w_j$ and thus we simply do nothing (p_i will never be returned by ACCESSMIN). For the latter case, we know that $f_{i_{j'}}(p_i - 1) > f_i(p_i - 1)$ and thus we simply set p_i to $p_i - 1$. We deal with the same case for q_i similarly.

- **DELETEMIN(\mathcal{H}):** We delete the node with minimum value from B (which is the one on top of the heap). This deleted node corresponds to one of the values $f_1(p_1), f_2(p_2), \dots, f_1(q_1), f_2(q_2), \dots$. Assume that $f_i(p_i)$ (resp. $f_i(q_i)$) is such value. We insert a node with value $f_i(p_i - 1)$ (resp. $f_i(q_i + 1)$).




Figure 3: Inserting a new function

2.3 Using the data structure for semi-matching problem

For any right vertex v , we construct a data structure H_v as in Section 2.2 to maintain f_{uv} for all neighbor of v which comes along with g_{uv} . These functions satisfy the properties above, as shown in Section 2.1. (We note that once x is deleted, $f_{*v}(x)$ will never change since this corresponds to adding a vertex v_x to the shortest path tree with distance $f_{*v}(x)$.)

The last issue is how to find γ_{uv} , the lowest point of an edge uv quickly. We now show an algorithm that finds γ_{uv} , for every edge $uv \in E$ in time $O(|V| + |E|)$ in total. This algorithm can be run before we start each iteration of the main algorithm (i.e., above Line 4 of Algorithm 2.1). To derive such algorithm, we need the following observation.

Lemma 2.6. *For any $v \in V$ and $u_1, u_2 \in U$, if $w_{u_1v} \geq w_{u_2v}$, then $\gamma_{u_1v} \leq \gamma_{u_2v}$.*

Proof. Note that by Lemma 2.5, γ_{uv} is the minimum integer $i \in [\deg(v)]$ such that $p(v^{i+1}) - p(v^i) \leq w_{uv}$. Also, for any $j < q(u_1v)$, $p(v^{j+1}) - p(v^j) > w_{u_1v}$ by definition. If $\gamma_{u_1v} > \gamma_{u_2v}$, then $p(v^{q(u_2v)+1}) - p(v^{q(u_2v)}) > w_{u_1v}$. However, $p(v^{q(u_2v)+1}) - p(v^{q(u_2v)}) \leq w_{u_2v}$. So, $w_{u_1v} < w_{u_2v}$. \square

Algorithm: The following algorithm finds γ_{uv} for all $uv \in E$. First, in the preprocessing step (which is done once before we begin the main algorithm), we order edges incident to v decreasingly by their weights, for every vertex $v \in V$. This process takes $O(\deg(v) \log(\deg(v)))$. We only have to compute γ_{uv} once, so this process does not affect the overall running time.

Next, for any $v \in V$, suppose that the list is $(u_1, u_2, \dots, u_{\deg(v)})$. Since $w_{u_1} \geq w_{u_2} \geq \dots \geq w_{u_{\deg(v)}}$, it implies that $\gamma_{u_1v} \leq \gamma_{u_2v} \leq \dots \gamma_{u_{\deg(v)}v}$ by Lemma 2.6. So, we first find γ_{u_1v} and then γ_{u_2v} and so on. This step takes $O(\deg(v))$ for each $v \in V$ and $O(m)$ in total. Therefore, the running time for computing the minimum point γ_{uv} 's is $O(m \log n)$.

3 Unweighted semi-matching

In this section, we present an algorithm that finds the optimal semi-matching in unweighted graph in $O(m\sqrt{n} \log n)$ time.

Overview

Our algorithm consists of the following three steps.

In the first step, we reduce the problem to the min-cost flow problem, using the same reduction from Harvey et al. [19]. (See Figure 4.) The details are provided in Section 3.1. We note that the flow is optimal if and only if there is no cost reducing path (to be defined later). We start with an arbitrary semi-matching and use this reduction to get a corresponding flow. The goal is to eliminate all the cost-reducing paths.

The second step is a divide-and-conquer algorithm used to eliminate all the cost-reducing paths. We call this algorithm CANCELALL (cf. Algorithm 3.1). The main idea here is to divide the graph into two subgraphs so that eliminating cost reducing paths “inside” each subgraph does not introduce any new cost reducing paths going through the other. This dividing step needs to be done carefully. We treat this in Section 3.2.

Finally, in the last component of the algorithm we deal with eliminating cost-reducing paths between two sets of vertices quickly. Naively, one can do this using any unit-capacity max-flow algorithm, but this does not give an improvement on the running time. To get a faster algorithm, we observe that the structure of the graph is similar to a *unit network*, where every vertex has in-degree or out-degree one. Thus, we get the same performance guarantee as the Dinic’s algorithm [7, 8].⁴ Details of this part can be found in Section 3.3.

After presenting the algorithm in the next three sections, we analyze the running time in Section 3.4. We note that this algorithm also works in a more general cost function (discussed in Section 3.5). We also observe an $O(n^{5/2} \log n)$ -time algorithm that arises directly from the reduction of the weighted case (discussed in Appendix B). This already gives an improvement over the previous results but our result presented here improve the running time further.

3.1 Reduction to min-cost flow and optimality characterization (revisited)

In this section, we review the characterization of the optimality of the semi-matching in the min-cost flow framework. We use the reduction as given in [19]. Given a bipartite graph $G = (U \cup V, E)$, we construct a directed graph N as follows. Let Δ denote the maximum degree of the vertices in V . First, add a set of vertices, called *cost centers*, $C = \{c_1, c_2, \dots, c_\Delta\}$ and connect each $v \in V$ to c_i with edges of capacity 1 and cost i , for all $1 \leq i \leq \deg(v)$. Second, add s and t as a source and sink vertex. For each vertex in U , add an edge from s to it with zero cost and unit capacity. For each cost center c_i , add an edge to t with zero cost and infinite capacity. Finally, direct each edge $e \in E$ from U to V with capacity 1 and cost 0. Observe that the new graph N has $O(n)$ vertices and $O(m)$ edges, and any semi-matching in G corresponds to a max flow in N .

Observe that the new graph N contains $O(n)$ vertices and $O(m)$ edges. It can be seen that any semi-matching in G corresponds to a max flow in N . (See example in Figure 4.) Moreover, Harvey et al. [19] prove that an optimal semi-matching in G corresponds to a min-cost flow in N ; in other words, the reduction described above is correct. Our algorithm based on observation that the largest cost is $O(|U|)$. This allows one to use the cost-scaling framework to solve the problem.

Now, we review an optimality characterization of the min-cost flow. We need to define a *cost-reducing path* first. Let R_f denote the residual graph of N with respect to a flow f . We call any path p from a cost center c_i to c_j in R_f an *admissible path* and call p a *cost-reducing path* if $i > j$.

⁴The algorithm is also known as “Dinic’s algorithm”. See [8] for details.




Figure 4: Reduction to the min-cost flow problem. Each edge is labelled with **(cost, capacity)** constraint. Thick edges either are matching edges or contain the flow.

A cost-reducing path is one-to-one corresponding to a negative cost cycle implying the condition for the minimality of f . Harvey *et al* [19] proved the following.

Lemma 3.1 ([19]). *A flow f is a min-cost flow in N if and only if there is no cost-reducing path in $R_f(N)$.*

Proof. Note that f is a min-cost flow if and only if there is no negative cycle in R_f . To prove the “only if” part, assume that there is an cost-reducing path from c_i to c_j . We consider the shortest one, i.e., no cost center is contained the path except the first and the last vertices. The edges that effect the cost of this path are only the first and the last ones because only edges incident to cost centers have cost. Cost of the first and the last edge is $-i$ and j respectively. Connecting c_i and c_j with t results a cycle of cost $j - i < 0$.

For the “if” part, assume that there is a negative-cost cycle in R_f . Consider the shortest cycle which contains only two cost centers, say c_i and c_j where $i > j$. This cycle contains an admissible path from c_i to c_j . \square

Given a max-flow f and a cost-reducing path P , one can find a flow f' with lower cost by augmenting f along P with a unit flow. This is later called *path cancelling*. We are now ready to explain our algorithm.

3.2 The divide-and-conquer algorithm

Our algorithm takes a bipartite graph $G = (U \cup V, E')$ and outputs the optimal semi-matching. It starts by transforming G into a graph N as described in the previous section. Since the source s and the sink t are always clear from the context, the graph N can be seen as a tripartite graph with vertices $U \cup V \cup C$; later on, we denote $N = (U \cup V \cup C, E)$. The algorithm proceeds by finding an arbitrary max-flow f from s to t in N which corresponds to a semi-matching in G . This can be done in linear time since the flow is equivalent to any semi-matching in G .

To find the min-cost flow in N , the algorithm uses a subroutine called CANCELALL (cf. Algorithm 3.1) to cancel all cost-reducing paths in f . Lemma 3.1 ensures that the final flow is optimal.

CANCELALL works by dividing C and solves the problem recursively. Given a set of cost centers C , the algorithm divides C into roughly equal-size subsets C_1 and C_2 such that, for any $c_i \in C_1$ and $c_j \in C_2$, $i < j$. This guarantees that there is no cost reducing path from C_1 to C_2 . Then it cancels all cost reducing paths from C_2 to C_1 by calling CANCEL algorithm (described in Section 3.3).

It is left to cancel the cost-reducing paths “inside” each of C_1 and C_2 . This is done by partitioning the vertices of N (except s and t) and forming two subgraphs N_1 and N_2 . Then solve the

Algorithm 3.1 CANCELALL($N = (U \cup V \cup C, E)$)

- 1: **if** $|C| = 1$ **then** halt **endif**
 - 2: Divide C into C_1 and C_2 of roughly equal size.
 - 3: CANCEL(N, C_2, C_1). {Cancel all cost-reducing paths from C_2 to C_1 }.
 - 4: Divide N into N_1 and N_2 where N_2 is “reachable” from C_2 and N_1 is the rest.
 - 5: Recursively solve CANCELALL(N_1) and CANCELALL(N_2).
-

problem separately on each of them. In more detail, we partition the graph N by letting N_2 be a subgraph induced by vertices reachable from C_2 in the residual graph and N_1 be the subgraph induced by the rest vertices. (Note that both graphs have s and t .) For example, in Figure 4, v_1 is reachable from c_3 by the path c_3, v_2, u_2, v_1 in the residual graph.

Lemma 3.2. CANCELALL(N) (cf. Algorithm 3.1) cancels all cost-reducing paths in N .

Proof. Recall that all cost-reducing paths from C_2 to C_1 are cancelled in line 3. Let S denote the set of vertices reachable from C_2 .

Claim 3.3. After line 3, no admissible paths between two cost centers in C_1 intersect S .

Proof. Assume, for the sake of contradiction, that there exists an admissible path from x to y , where $x, y \in C_1$, that contains a vertex $s \in S$. Since s is reachable from some vertex $z \in C_2$, there must exist an admissible path from some vertex in z to y ; this leads to a contradiction. \square

This claim implies that, in our dividing step, all cost-reducing paths between pairs of cost centers in C_1 remain entirely in N_1 . Furthermore, vertices in any cost reducing path between pairs of cost centers in C_2 must be reachable from C_2 ; thus, they must be inside S . Therefore, after the recursive calls, no cost-reducing paths between pairs of cost centers in the same subproblems C_i are left. The lemma follows if we can show that in these processes we do not introduce more cost-reducing paths from C_2 to C_1 . To see this, note that all edges between N_1 and N_2 remain untouched in the recursive calls. Moreover, these edges are directed from N_1 to N_2 , because of the maximality of S . Therefore there is no admissible path from C_2 to C_1 . \square

3.3 Cancelling paths from C_2 to C_1

In this section we describe an algorithm that cancels all admissible paths from C_2 to C_1 in R_f , which can be done by finding a max flow from C_2 to C_1 . To simplify the presentation, we assume that there is a super-source s and super-sink t connecting to vertices in C_2 and in C_1 , respectively.

To find a maximum flow, observe that N is unit-capacity and every vertex of U has indegree 1 in R_f . By exploiting these properties, we show that Dinitz’s blocking flow algorithm [7] can find a maximum flow in $O(|E|\sqrt{|U|})$ time. The algorithm is done by repeatedly augmenting flows through the shortest augmenting paths. (see Appendix C).

Lemma 3.4. Let d_i be the length of the shortest $s - t$ path in the residual graph at the i^{th} iteration. For all i , $d_{i+1} > d_i$.

The lemma can be used to show that Dinitz’s algorithm terminates after n rounds of the blocking flow step, where n is the number of vertices. Since after the n -th round, the distance between the source is more than n , which means that there is no augmenting path from s to t

in the residual graph. The number of rounds can be improved for certain classes of problems. Even and Tarjan [10] and Karzanov [22] showed that in unit capacity networks, Dinitz's algorithm terminates after $\min(n^{2/3}, m^{1/2})$ rounds, where m is the number of edges. Also, in unit networks, where every vertex has in-degree one or out-degree one, Dinitz's algorithm terminates in $O(\sqrt{n})$ time (see, e.g., Tarjan's book [46]). Since the graph N we are considering is very similar to unit networks, we are able to show that Dinitz's algorithm also terminates in $O(\sqrt{n})$ in our case.

For any flow f , a *residual flow* f' is a flow in a residual graph R_f of f . If f' is maximum in R_f , $f + f'$ is maximum in the original graph. The following lemma relates the amount of the maximum residual flow with the shortest distance from s to t in our case. The proof is a modification of Theorem 8.8 in [46].

Lemma 3.5. *If the shortest $s - t$ distance in the residual graph is $d > 4$, the amount of the maximum residual flow is at most $O(|U|/d)$.*

Proof. A maximum residual flow in a unit capacity network can be decomposed into a set \mathcal{P} of edge-disjoint paths where the number of paths equals to the flow value. Each of these paths are of length at least d . Clearly, each path contains the source, the sink, and exactly two cost centers. Now consider any path $P \in \mathcal{P}$ of length l . It contains $l - 3$ vertices from $U \cup V$. Since the original graph is a bipartite graph, at least $\lfloor (l - 3)/2 \rfloor \geq \lfloor (d - 3)/2 \rfloor \geq (d - 4)/2$ vertices are from U . Note that each path in \mathcal{P} contains a disjoint set of vertices in U , since a vertex in U has in-degree one. Therefore, we conclude that there are at most $2|U|/(d - 4)$ paths in \mathcal{P} . The lemma follows since each path has one unit of flows. \square

From these two lemma, we have the main lemma for this section.

Lemma 3.6. *CANCEL terminates in $O(|E|\sqrt{|U|})$ time.*

Proof. Since each iteration can be done in $O(|E|)$ time, it is enough to prove that the algorithm terminates in $O(\sqrt{|U|})$ rounds. The previous lemma implies that the amount of the maximum residual flow after the $O(\sqrt{|U|})$ -th rounds is $O(\sqrt{|U|})$ units. The lemma thus follows because after that the algorithm augments at least one unit of flow for each round. \square

3.4 Running time

The running time of the algorithm is dominated by the running time of CANCELALL, which can be analyzed as follows. Let $T(n, n', m, k)$ denote the running time of the algorithm when $|U| = n$, $|V| = n'$, $|E| = m$, and $|C| = k$. For simplicity, assume that k is a power of two. By Lemma 3.6, CANCEL runs in $O(|E|\sqrt{|U|})$ time. Therefore,

$$T(n, n', m, k) \leq c \cdot m\sqrt{n} + T(n_1, n'_1, m_1, k/2) + T(n_2, n'_2, m_2, k/2),$$

for some constant c , where n_i, n'_i , and m_i denote the number of vertices and edges in N_i , respectively. Recall that each edge participates in at most one of the subproblems; thus, $m_1 + m_2 \leq m$. Observe that the number of cost centers always decrease by a factor of two. Thus, the recurrence is solved to $O(\sqrt{nm} \log k)$. Since $k = O(|U|)$, the running time is $O(\sqrt{nm} \log n)$ as claimed. Furthermore, the algorithm can work in more general cost function with the same running time as shown in the next section.

3.5 Generalizations of an unweighted algorithm

The problem can be viewed in a slightly more general version. In Harvey et al. [19], the cost functions for each vertex $v \in V$ are the same. We relax this condition, allowing different function for each vertex where each function is convex. More precisely, for each $v \in V$, let $f_v : \mathbb{Z}_+ \rightarrow \mathbb{R}$ be a convex function, i.e., for any i , $f_v(i+1) - f_v(i) \geq f_v(i) - f_v(i-1)$. The cost for matching M on vertex v is $f_v(\deg_M(v))$. In this convex cost function, the transformation similar to what described in Section 3.1 can still be done. However, the number of different values of f_v is now $O(|E|)$. So, the size of the set of cost centers C is now upper bounded by $O(|E|)$ not $O(|U|)$. Therefore, the running time of our algorithm becomes $O(|E|\sqrt{|U|}\log|C|) = O(|E|\sqrt{|U|}\log|E|) = O(\sqrt{nm}\log n)$ (since $|E| \leq n^2$) which is the same as before.

4 Extension to Balanced Edge Cover problem

Recall that the problem is the following.

Input: A simple undirected graph $G = (V, E)$.

Task: Find an *edge cover* F minimizing $c(F) = \sum_{v \in V} \deg_F(v)$, where $\deg_F(v) = |\{vu \in F\}|$ and we say that F is an edge cover if $\deg_F(v) \geq 1$ for all $v \in V$.⁵

We call the solution of such problem an *optimal balanced edge cover*. Observe that any minimal edge cover – including any optimal balanced edge cover – induces a star forest; i.e., every connected component has at most one vertex of degree greater than one (we call such vertices centers) and the rest have degree exactly one. For any optimal balanced edge cover F , we call any set of vertices C an *extended set of centers of F* if C contains all centers of F and exactly one vertex from each of the connected components that have no center. (To be precise, C is a center if C contains all centers of F and each connected component in the subgraph induced by F contains exactly one vertex in C .)

To solve the balanced edge cover problem using semi-matching algorithm, we first make a further observation that if an extended set of centers is given, then optimal balanced edge cover can be found by simply solving the unweighted semi-matching problem.

Lemma 4.1. *Let C be an extended set of centers of some optimal balanced edge cover F . Let $G' = ((V \setminus C) \cup C, E')$ be a bipartite graph where E' is the set of edges between $V \setminus C$ and C in G . Then, an optimal semi-matching in G' (where we allow vertices in C to be connected more than once) is an optimal balanced edge cover in G .*

Proof. Let M be any optimal semi-matching in G' . First, note that F is also a semi-matching in G' . Thus, the cost of M is less than the cost of F . It is left to show that M is an edge cover. In other words, it is left to prove that every vertex in C is covered by M .

Assume for the sake of contradiction that there is a vertex $v \in C$ that is not covered by M . We show that there exists a cost-reducing path of M starting from v as follows. Starting from

⁵We note that the original definition of the balanced edge cover problem has a function $f : \mathbb{Z}^+ \rightarrow \mathbb{R}^+$ as an input [16]. However, it is shown in [16] that the optimal balanced edge cover can be determined independently of function f as long as f is strictly monotonic. In other words, the problem is equivalent to the one we define here.

$v_0 = v$, let v_1 be any vertex adjacent to v_0 in F and v_2 be a (unique) vertex adjacent to v_1 in M . If $\deg_M(v_2) > 1$, then we stop the process. Otherwise, repeat the process by finding a vertex v_3 adjacent to v_2 in F and a vertex v_4 adjacent to v_3 . Observe that all vertices found during the process are unique since every vertex in $V \setminus C$ has degree exactly one in F , every vertex found in the process, except v_0 , has degree exactly one in M , and there is no edge in F between two vertices in C . Therefore, the process above must stop. Moreover, the path obtained after the process stops is a cost-reducing path, contradicting the assumption that M is an optimal semi-matching. \square

It is left to find any extended set of centers. We do so by define levels of vertices based on some edge cover.

Definition 4.2. For any edge cover F , define the *levelling* of vertices in F , denoted by L_F , as follows.

First, let all center vertices (i.e., all vertices with degree more than one in F) be on level 1. For $i = 1, 2, \dots$, we construct level $i + 1$ by looking at any vertex v not yet assigned to any level. If i is odd and v shares an edge in F with a vertex on level i , then we add v to level $i + 1$. Otherwise, we add v to level $i + 1$ if i is even, v shares an edge not in F with a vertex on level i and v does not share an edge in F with any vertex on level $i + 1$. (For example, after we put vertices of degree more than one in the first level, we put their leaves on level two. Then, we put all vertices adjacent to these leaves by non-covering edges on level three. But, if we see a single edge having both end vertices on level three, we put one end to level four and so on.)

Note that when the process finishes, there might be some vertices that are not assigned to any level. \square

An optimal balanced edge cover can be found by the following algorithm.

FIND-CENTER Algorithm: First, find a minimum cardinality edge cover F . Then, find L_F in a breadth-first manner. Let M be an optimal semi-matching of a bipartite graph where the left vertices are even-level vertices and right vertices are odd-level vertices. Output M and edges in F between vertices with no level.

Now, we show the running time and the correctness of FIND-CENTER algorithm.

Running time analysis: F can be found by simply adding uncovered vertices to a maximum cardinality matching [13, 35]. The maximum cardinality matching in bipartite graph can be found by Micali-Vazirani algorithm [34] in $O(\sqrt{nm})$, or in $O(n^\omega)$ by Harvey algorithm [18], where ω is a time for computing matrix multiplication. However, since the running time of semi-matching algorithm is $O(\sqrt{nm} \log n)$. It suffices to use the first one. Thus, F can be found in $O(\sqrt{nm})$ time. Moreover, finding L_F could be done in a breadth-first manner which takes $O(n + |F| + |M|) = O(n)$ time. Therefore, the time for the reduction from balanced edge cover to semi-matching problem is $O(\sqrt{nm})$ implying the total running time of $O(\sqrt{nm} \log n)$.

Correctness:

The proof of correctness uses an algorithm BEC1 proposed in [16]. This algorithm starts from any minimum edge cover and keep augmenting along a *cost-reducing* path until such path does not exist. Here a cost-reducing path regarding to an edge cover F is a path starting from any center vertex u , follow any edge in F and follow an edge not in F . The path keeps using edges in F and

edges not in F alternately until it finally uses an edge not in F and ends at a vertex v such that $\deg_F(v) \leq \deg_F(u) - 2$. (See [16] for the formal definition.) It is shown that BEC1 returns an optimal balanced edge cover.

Lemma 4.3. *Let C be the set returned from the FIND-CENTER algorithm. Then C is an extended set of centers of some optimal balanced edge cover F^* . In other words, there exists F^* such that all of its centers are in C and each connected component (in the subgraph induced by F^*) has exactly one vertex in C .*

Proof. Let F be the minimum cardinality edge cover found by the FIND-CENTER algorithm. Consider a variation of BEC1 algorithm where we augment along a shortest cost-reducing path. We claim that we can always augment along the shortest cost-reducing path in such a way that parity of vertices' levels never change. To be precise, we construct a sequence of minimum cardinality edge covers $F = F_1, F_2, \dots$ where we get F_i from F_{i-1} by augmenting along some shortest cost-reducing path. By the following process, we guarantee that if any vertex is on an odd (even) level in L_F then it is on an odd (even) level in L_{F_i} . Moreover, if a vertex belongs to no level in L_F then it has no level in L_{F_i} .

Suppose we are at F_i and our guarantee is maintained so far. Let P be any shortest cost-reducing path on F_i . If there is no such P , then we found F^* . Otherwise, we consider two cases.

Case 1 If P contains only vertices on level 1 and 2: This is equivalent to reconnecting vertices on level 2 to vertices on level 1. Level of every vertex is the same in L_{F_i} and $L_{F_{i+1}}$. Thus, the guarantee is maintained.

Case 2 Otherwise: Let $P = v_0 v_1 v_2 \dots v_k$. Recall that k is even and note that all of v_0, v_1, \dots, v_{k-1} must be on level 1 and 2 (alternately); otherwise, we can stop at the first vertex that we visit on other level and obtain a shorter cost-reducing path. Now, let us augment from v_0 until we reach v_{k-2} . At this point, v_{k-1} must have degree at least three (after augmentation) because it is on level 1 (which means that it has degree more than one in F_i) and just receives one more edge from augmentation. If v_k is on level 3, then we are done as it will be on level 1 in $L_{F_{i+1}}$ and all vertices in its subtree will be 2 levels higher. If not, then v_k must be on level 4. Let a be a vertex adjacent to v_k by an edge in F_i (which is on level 3) and let b be a vertex on level 2 adjacent to a (by an edge not in F_i). There are two subcases.

Case 2.1 When $v_{k-1} = b$: In this case, we use a path $v_1 v_2 \dots v_{k-1} a$ instead.

Case 2.2 When $v_{k-1} \neq b$: In this case, we get an edge cover with cardinality smaller than $|F_i| = |F|$ by deleting three edges in F_i incident to b , v_{k-1} and v_k and add ab and $v_{k-1} v_k$. (Note that for the case that b is covered by an edge incident to v_{k-2} , we use the fact that v_{k-2} has degree at least 3 noted earlier.) So, this case is impossible as it contradicts the fact that F is minimum cardinality edge cover. \square

Acknowledgment We thank David Pritchard for useful suggestions, Jane (Pu) Gao for pointing out some related surveys, and Dijun Luo for pointing out some errors in the earlier version of this paper.

References

- [1] D. Abraham. Algorithmics of two-sided matching problems. Master's thesis, Department of Computer Science, University of Glasgow, 2003.
- [2] Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network flows: theory, algorithms, and applications*. Prentice-Hall, Inc., 1993.
- [3] J. Blažewicz, K. Ecker, E. Pesch, G. Schmidt, and J. Weglarz. Handbook on scheduling: from theory to applications.
- [4] D. Bokal, B. Bresar, and J. Jerebic. A generalization of hungarian method and hall's theorem with applications in wireless sensor networks. *arXiv/0911.1269*, 2009.
- [5] John L. Bruno, Edward G. Coffman Jr., and Ravi Sethi. Algorithms for minimizing mean flow time. In *IFIP Congress*, pages 504–510, 1974.
- [6] John L. Bruno, Edward G. Coffman Jr., and Ravi Sethi. Scheduling independent tasks to reduce mean finishing time. *Commun. ACM*, 17(7):382–387, 1974.
- [7] E. A. Dinic. Algorithm for solution of a problem of maximum flow in networks with power estimation. *Soviet Mathematics Doklady*, 11:1277–1280, 1970.
- [8] Yefim Dinitz. Dinitz' algorithm: The original version and even's version. In *Essays in Memory of Shimon Even*, pages 218–240, 2006.
- [9] J. Edmonds and R.M. Karp. Theoretical improvements in algorithmic efficiency for network flow problems. In *Combinatorial Structures and Their Applications*, pages 93–96. Gordon and Breach, New York, 1970.
- [10] Shimon Even and R. Endre Tarjan. Network flow and testing graph connectivity. *SIAM J. Comput.*, 4(4):507–518, December 1975.
- [11] Jittat Fakcharoenphol, Bundit Laekhanukit, and Danupon Nanongkai. Faster algorithms for semi-matching problems (extended abstract). In *ICALP (1)*, pages 176–187, 2010.
- [12] Harold N. Gabow and Robert Endre Tarjan. Faster scaling algorithms for network problems. *SIAM J. Comput.*, 18(5):1013–1036, 1989.
- [13] T. Gallai. Über extreme Punkt-und Kantenmengen. *Ann. Univ. Sci. Budapest, Eotvos Sect. Math.*, 2:133–138, 1959.
- [14] R.L. Graham, E.L. Lawler, J.K. Lenstra, and A.H.G.R. Kan. Optimization and approximation in deterministic sequencing and scheduling: a survey. *Annals of Discrete Mathematics*, 5(2):287–326, 1979.
- [15] Yuta Harada, Hirotaka Ono, Kunihiro Sadakane, and Masafumi Yamashita. Optimal balanced semi-matchings for weighted bipartite graphs. *IPSJ Digital Courier*, 3:693–702, 2007.
- [16] Yuta Harada, Hirotaka Ono, Kunihiro Sadakane, and Masafumi Yamashita. The balanced edge cover problem. In *ISAAC*, pages 246–257, 2008.
- [17] Nicholas J. A. Harvey. Semi-matchings for bipartite graphs and load balancing (slide). <http://people.csail.mit.edu/nickh/Publications/SemiMatching/Semi-Matching.ppt>, July 2003.
- [18] Nicholas J. A. Harvey. Algebraic algorithms for matching and matroid problems. *SIAM J. Comput.*, 39(2):679–702, 2009. Also appeared in FOCS'06.
- [19] Nicholas J. A. Harvey, Richard E. Ladner, László Lovász, and Tami Tamir. Semi-matchings for bipartite graphs and load balancing. *J. Algorithms*, 59(1):53–78, 2006. Conference version in WADS'03.
- [20] WA Horn. Minimizing average flow time with parallel machines. *Operations Research*, pages 846–847, 1973.
- [21] Ming-Yang Kao, Tak Wah Lam, Wing-Kin Sung, and Hing-Fung Ting. An even faster and more unifying algorithm for comparing trees via unbalanced bipartite matchings. *J. Algorithms*, 40(2):212–233, 2001.
- [22] A.V. Karzanov. On finding maximum flows in networks with special structure and some applications. *Matematicheskie Voprosy Upravleniya Proizvodstvom*, 5:81–94, 1973.
- [23] Anshul Kothari, Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Congestion games, load balancing, and price of anarchy. In *CAAN*, pages 13–27, 2004.
- [24] K. Lee, J. Y-T. Leung, and M. L. Pinedo. Scheduling jobs with equal processing times subject to

- machine eligibility constraints. *Working Paper*, 2009.
- [25] Kangbok Lee, Joseph Y. T. Leung, and Michael L. Pinedo. A note on “an approximation algorithm for the load-balanced semi-matching problem in weighted bipartite graphs”. *Inf. Process. Lett.*, 109(12):608–610, 2009.
 - [26] Joseph Y.-T. Leung and Chung-Lun Li. Scheduling with processing set restrictions: A survey. *International Journal of Production Economics*, 116(2):251–262, December 2008.
 - [27] Chung-Lun Li. Scheduling unit-length jobs with machine eligibility restrictions. *European Journal of Operational Research*, 174(2):1325–1328, October 2006.
 - [28] Yixun Lin and Wenhua Li. Parallel machine scheduling of machine-dependent jobs with unit-length. *European Journal of Operational Research*, 156(1):261–266, July 2004.
 - [29] R. Logendran and F. Subur. Unrelated parallel machine scheduling with job splitting. *IIE Transactions*, 36(4):359–372, 2004.
 - [30] László Lovász. The membership problem in jump systems. *J. Comb. Theory, Ser. B*, 70(1):45–66, 1997.
 - [31] Chor Ping Low. An efficient retrieval selection algorithm for video servers with random duplicated assignment storage technique. *Inf. Process. Lett.*, 83(6):315–321, 2002.
 - [32] Chor Ping Low. An approximation algorithm for the load-balanced semi-matching problem in weighted bipartite graphs. *Inf. Process. Lett.*, 100(4):154–161, 2006. Also appeared in TAMC 2006.
 - [33] Renita Machado and Sirin Tekinay. A survey of game-theoretic approaches in wireless sensor networks. *Computer Networks*, 52(16):3047–3061, 2008.
 - [34] Silvio Micali and Vijay V. Vazirani. An $o(\sqrt{|V|}|E|)$ algorithm for finding maximum matching in general graphs. In *FOCS*, pages 17–27, 1980.
 - [35] R.Z. Norman and M.O. Rabin. An algorithm for a minimum cover of a graph. *Proceedings of the American Mathematical Society*, pages 315–319, 1959.
 - [36] Michael Pinedo. *Scheduling: Theory, Algorithms, and Systems (2nd Edition)*. Prentice Hall, August 2001.
 - [37] Narayanan Sadagopan, Mitali Singh, and Bhaskar Krishnamachari. Decentralized utility-based sensor network design. *Mob. Netw. Appl.*, 11(3):341–350, 2006.
 - [38] Alexander Schrijver. *Combinatorial optimization : polyhedra and efficiency*. volume A, paths, flows, matchings, chapter 1-38. Springer, 2003.
 - [39] René Sitters. Two np-hardness results for preemptive minsum scheduling of unrelated parallel machines. In *IPCO*, pages 396–405, 2001.
 - [40] C. D. Spyropoulos and D. J. Evans. Analysis of the q.a.d. algorithm for an homogeneous multiprocessor computing model with independent memories. *International Journal of Computer Mathematics*, pages 237–255, 1985.
 - [41] L.H. Su. Scheduling on identical parallel machines to minimize total completion time with deadline and machine eligibility constraints. *The International Journal of Advanced Manufacturing Technology*, 40(5):572–581, 2009.
 - [42] Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Uncoordinated load balancing and congestion games in p2p systems. In *IPTPS*, pages 123–130, 2004.
 - [43] Subhash Suri, Csaba D. Tóth, and Yunhong Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007. Also appeared in SPAA’04.
 - [44] A. Tamir. Least majorized elements and generalized polymatroids. *Mathematics of Operations Research*, pages 583–589, 1995.
 - [45] T. Tamir and B. Vaksendiser. Algorithms for storage allocation based on client preferences. *International Symposium on Combinatorial Optimization*, 2008.
 - [46] Robert Endre Tarjan. *Data structures and network algorithms*. Society for Industrial and Applied Mathematics, 1983.
 - [47] N. Tomizawa. On some techniques useful for solution of transportation network problems. In *Networks 1*, pages 173–194, 1971.
 - [48] Z. Vaik. On scheduling problems with parallel multi-purpose machines. Technical report, Technical Reports, Egervary Research Group on Combinatorial Optimization, Hungary, [www. cs. elte](http://www.cs.elte).

APPENDIX

A Edmonds-Karp-Tomizawa algorithm for weighted bipartite matching

In this section, we briefly explain Edmonds-Karp-Tomizawa (EKT) algorithm. The algorithm starts with an empty matching M and iteratively augments (i.e., increases the size of) M . The matching in each iteration is maintained so that it is *extreme*; i.e., it has highest weight among matching of the same cardinality. The augmenting procedure is as follows. Let M be a matching maintained so far. Let D_M be the directed graph obtained from \hat{G} by orienting each edge e in M from \hat{V} to U with length $\ell_e = -w_e$ and orienting each edge e not in M from U to \hat{V} with length $\ell_e = w_e$. Let U_M (respectively, \hat{V}_M) be the set of vertices in U (respectively, \hat{V}) not covered by M . If $|M| \neq |U|$, then there is a U_M - \hat{V}_M path. Find a shortest such path, say P , and augment M along P ; i.e., set $M = M \Delta P$. Repeat with the new value of M until $|M| = |U|$.

The bottleneck of this algorithm is the shortest path algorithm. Although D_M has negative edge length, one can find a *potential* and applying Dijkstra's algorithm on a graph D_M with non-negative *reduced cost*. The potential and reduced cost are defined as follows.

Definition A.1. A function $p : U \cup \hat{V} \rightarrow \mathbb{R}$ is a *potential* if, for every edge uv in the residual graph D_M , $\tilde{\ell}_{uv} = \ell_{uv} + p(u) - p(v)$ is non-negative. We call $\tilde{\ell}$ a *reduced cost* with respect to a potential p .

The key idea of using a potential is that a shortest path from u to v with respect to a reduced cost $\tilde{\ell}$ is also a shortest with respect to ℓ . We omit details here (see, e.g., ([38, Chapter 7 and Section 17.2]), but note that we can use a distance function found in the last iteration of the algorithm as a potential, as in Algorithm 2.1.

Dijkstra's algorithm.

We now explain Dijkstra's algorithm on graph D_M with non-negative edge weight defined by $\tilde{\ell}$. Our presentation is slightly different from the standard one but will be easy to modify later. The algorithm keeps a subset X of $U \cup \hat{V}$, called the set of undiscovered vertices, and a function $d : U \cup \hat{V} \rightarrow \mathbb{R}^+$ (the *tentative distance*). Start with $X = U \cup \hat{V}$ and set $d(u) = 0$ for all $u \in U_M$ and $d(v) = \infty$ for all vertex $v \notin U_M$. Apply the following iteratively:

- 1: Find $u \in X$ minimizing $d(u)$ over $u \in X$. Set $X = X \setminus \{u\}$.
- 2: For each neighbor v of u in D_M , “relax” uv : set $d(v) \leftarrow \min\{d(v), d(u) + \tilde{\ell}_{uv}\}$.

The running time of Dijkstra's algorithm depends on the implementation. One implementation is by using Fibonacci heap. Each vertex $v \in U \cup \hat{V}$ is kept in the heap with key $d(v)$. Finding and extracting a vertex of minimum tentative distance can be done in an amortized time bound of $O(\log |U \cup \hat{V}|)$ by “extract-min”, and relaxing an edge can be done in an amortized time bound of $O(1)$ by “decrease-key”.

Consider the running time caused by finding a shortest path. Let $n = |U \cup V|$ and $m = |E|$. Then we have to call insertion $O(n)$ times, decrease-key $O(m)$ times, and extract-min $O(n)$ times. Thus, the overall running time is $O(m + n \log n)$.

B Observation: $O(n^3)$ and $O(n^{5/2} \log(nW))$ time algorithms

Recall to the reduction from the weighted semi-matching problem to the weighted bipartite matching problem, or equivalently, an assignment problem. The reduction was shown in [5, 20, 17]. We include it here for completeness. Given a bipartite graph $G = (U \cup V, E)$ with edge weight w , an instance for the semi-matching problem, we construct a bipartite graph $\hat{G} = (U \cup \hat{V}, \hat{E})$ with weight \hat{w} , an instance for the weighted bipartite matching problem, as follows. For every vertex $v \in V$ of degree $\deg(v)$, we create *exploded* vertices $v^1, v^2, \dots, v^{\deg(v)}$ in \hat{V} and let \hat{V}_v denote a set of such vertices. For each edge uv in E of weight w_{uv} , we also create $\deg(v)$ edges $uv^1, uv^2, \dots, uv^{\deg(v)}$, with associated weights $w_{uv}, 2 \cdot w_{uv}, \dots, \deg(v) \cdot w_{uv}$, respectively. It is easy to verify that finding optimal semi-matching in G is equivalent to finding minimum matching in \hat{G} . Figure 1(a) shows an example of this reduction.

The construction yields a graph \hat{G} with $O(m)$ vertices and $O(nm)$ edges. Applying any existing algorithms for weighted bipartite matching directly is not enough to get an improvement. However, we observe that the reduction can be done in $O(n^2 \log n)$ time, and we can apply the result of Kao et al. in [21] to reduce the number of participating edges to $O(n^3)$. Thus, Gabow and Tarjan's scaling algorithm [12] give us the following result.

Observation B.1. *If all edges have non-negative integer weight bounded by W , then there is an algorithm for the weighted semi-matching problem with the running time of $O(n^{5/2} \log nW)$.*

This result immediately gives an $O(n^{5/2} \log n)$ time algorithm for the unweighted case (i.e., $W = 1$). Hence, we already have an improvement upon the previous $O(nm)$ time algorithm for the case of dense graph.

Now, we give an explanation on the observation. If we reduce the problem normally (as in Section 2) to get \hat{G} , then the number of edges in \hat{G} and the running time will be $O(nm)$. However, since the size of any matching in the graph \hat{G} is at most $|U|$, it suffices to consider only the smallest $|U|$ edges in \hat{G} incident to each vertex in U . Therefore, we may assume that \hat{G} has $O(n^2)$ edges. (The same observation is also used in [21].)

More precisely, let E_u be a set of edges incident to u in \hat{G} , and R be a set of $|U|$ smallest edges of E_u . If the maximum matching of the minimum weight, say M , contains an edge $e \in E_u \setminus R$, then $U \cup \{e\}$ has $|U| + 1$ edges implying that there is an edge $e' \in R$ incident to a vertex $v \in \hat{V}$ not matched by M . Thus, we can replace e with e' which results in a matching of smaller weight. Therefore, we need to keep only $|U|^2$ edges in our reduction.

Moreover, we can also reduce the time for a reduction to $O(n^2 \log n)$ as well. The faster reduction can be as follows. For each vertex $u \in U$, we first add all edges incident to u , to a binary heap H with addition information $i = 1$, say (e, i) . Then we iteratively extract minimum $(e = uv, i)$ from H , create an edge uv^i in E' with weight $i \cdot w(e)$, and insert $(uv^{i+1}, i + 1)$ back to H . We repeat the process until u has $|U|$ incident edges in E' . The pseudocode of the reduction is given in Algorithm B.1.

Consider a vertex $u \in U$. In any time during the reduction, there are $O(\deg_{\hat{G}}(u))$ edges in H . So, the extract-min takes $O(\log(\deg_{\hat{G}}(u)))$ time. The time for inserting a vertex to \hat{V} and an edge to \hat{E} is $O(1)$ which is dominated by the time for extract-min. Thus, we have to consider only the time for heap operations. For each vertices $u \in U$, we have to call insertion $\deg_{\hat{G}}(u) + |U|$ times and extract-min $|U|$ times. Thus, the time required to process each vertex of U is $O((\deg_{\hat{G}}(u) + |U|) \log |U|)$. It follows that the total running time of the reduction is $O((|E| + |U|^2) \log |U|) = O(n^2 \log n)$.

Algorithm B.1 REDUCTION ($G = (U \cup V, E), w$)

```
1: Create an empty set  $\hat{E}, \hat{V}$ .
2: for all vertices  $u \in U$  do
3:   Create a binary heap  $H$ .
4:   for all edges  $e$  incident to  $u$  do
5:     Insert  $(e, 1)$  to  $H$  with key  $w(u)$ .
6:   end for
7:   for  $k \leftarrow 1$  to  $|U|$  do
8:     Extract-min from  $H$ , resulted in  $(e = uv, i)$ .
9:     Insert a vertex  $v^i$  to  $\hat{V}$  (If not exists).
10:    Insert an edge  $uv^i$  to  $\hat{E}$ .
11:    Insert an edge  $(e' = uv^{i+1}, i + 1)$  to  $H$  with key  $w(e) \cdot (i + 1)$ .
12:   end for
13:   Destroy a binary heap  $H$ .
14: end for
15: Return  $\hat{G} = (U \cup \hat{V}, \hat{E})$ .
```

Now, we run algorithms for bipartite matching problem on the graph \hat{G} with n^2 edges. Using Edmonds-Karp-Tomizawa, the running time becomes $O(nm) = O(n^3)$. Using Gabow-Tarjan's scaling algorithm, the running time becomes $O(\sqrt{n}m \log(nW) = O(n^{5/2} \log(nW)))$, where W is the maximum edge weight.

C Dinitz's blocking flow algorithm

In this section, we will give an outline of Dinitz's blocking flow algorithm [7]. Given a network R with source s and sink t , a flow g is a *blocking flow* in R if every path from the source to the sink contains a saturated edge, an edge with zero residual capacity. A blocking flow is usually called a greedy flow, since the flow cannot be increased without any rerouting of the previous flow paths. In a unit capacity network, depth-first search can be used to find blocking flow in linear time.

Dinitz's algorithm works in *layer graph*, a subgraph whose edges are in at least one shortest path from s to t . This condition implies that we only augment along the shortest paths. The algorithm proceeds by successively find blocking flows in the layer graphs of the residual graph of the previous round. The following is an important property (see, e.g., [2, 38, 46] for proofs). It states that the distance between the source and the sink always increases after each blocking flow step.

In the case of unit-capacity, Even-Tarjan [10] and Karzanov [22] showed that the algorithm finds a maximum flow in time $O(\min\{n^{2/3}, m^{1/2}\}m)$. In the case of *unit-network*, i.e., every vertex either has indegree 1 or outdegree 1, the algorithm finds a maximum flow in time $O(\sqrt{nm})$.